Dynamic Profile Selection & Cooperative Transcoding: An Efficient Future for the Multiscreen World

Santhana Chari, Ph.D., VP Engineering Digital Video Systems

Introduction

The advent of Adaptive Bit Rate (ABR) delivery has unlocked new doors as service providers move their content beyond the TV and onto the screens of PCs, tablets and smartphones. But while it is enabling new multiscreen video opportunities for service providers, this technology is creating new challenges, too. With the need to support so much content, so many devices, and a wide range of network conditions, transcoding requirements and their capital and operational costs are exploding. In addition, as customer adoption rises and the volume of multiscreen traffic increases, the burden on access networks is expected to rise.

This is causing solution vendors and service providers to examine the transcoding process and search for new ways of preparing video for its final destination in the multiscreen world. This paper focuses on two innovations – dynamic profile selection and cooperative transcoding – that have the potential to significantly reduce the processing and bandwidth requirements for more efficient multiscreen video delivery.

New Opportunities for Multiscreen Transcoding Efficiency

ABR transcoding begins by identifying the various screen resolutions, streaming protocols, conditional access methods and network conditions that need to be supported, and creating a number of unique profiles based on these attributes. When a consumer device such as a tablet or smartphone requests a piece of content, the correct file for the device and network conditions is either selected from storage or created by the transcoder, then streamed to the user. If the network becomes congested, a lower bit rate version is streamed instead to enable the video to play uninterrupted. Should bandwidth become abundant, a high bit rate version takes over to produce the highest possible video quality.

While it may seem simple, the transcoding resources required to make this process work smoothly are substantial. When we consider the sheer number of screen sizes for smart phones through HDTVs, add support for the most popular streaming profiles such as Apple HLS, Microsoft Smooth Streaming, and MPEG DASH, and factor in conditional access methods including PlayReady, SecureMedia and Adobe Access, today's transcoders are required to support numerous unique file versions for a given piece of content. And adding support for multiple bit rates adds another set of requirements to the already intensive transcoding workload. Meanwhile, increasing multiscreen adoption brings with it an increased bandwidth requirement that has the potential to strain service providers' access networks over time.

By harnessing dynamic profile selection and cooperative transcoding, service providers can attain significant networking and processing efficiencies to reduce expenses without sacrificing video quality.

Reducing Over- & Under- Processing with Dynamic Profile Selection

Today, video profiles are chosen ahead of time to correspond with the specific display devices service providers expect their subscribers to use when viewing multiscreen content. The profiles are selected not only to fit the needs of the supported devices, but also to achieve a targeted number of bits per pixel that balance video quality with impact on the network. While new profiles can be added or taken away to keep pace with device trends, the profiles are hardwired and not changed dynamically.

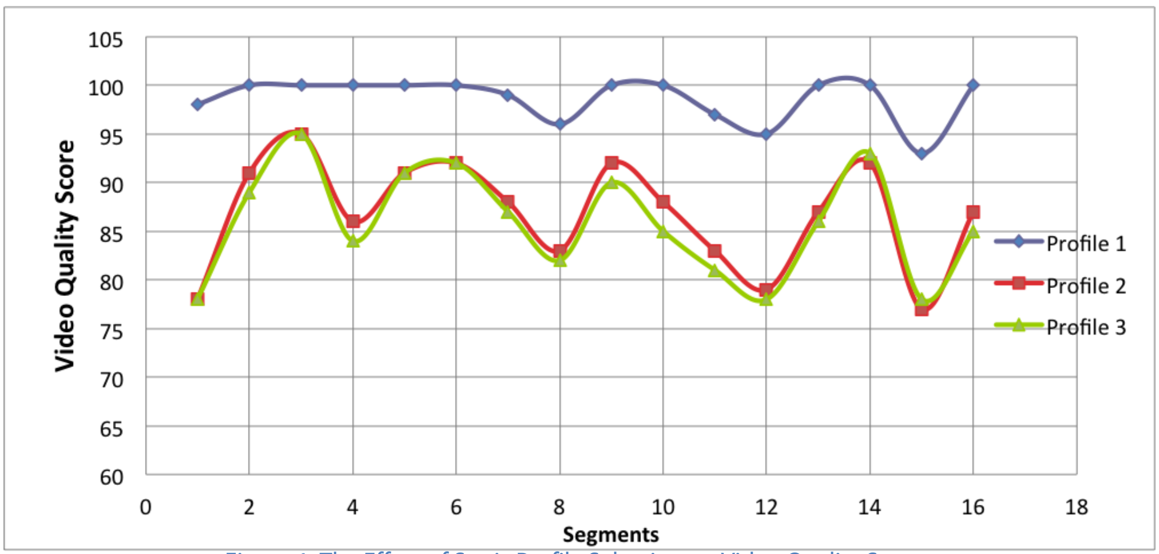

Figure 1: The Effect of Static Profile Selection on Video Quality Scores

However, due to the variability of video content, these handpicked static profiles often do not produce versions of content that are at the desired video quality, thereby not meeting the end user quality of experience expectation, or placing an undue burden on the transcoder. An example of over-processing is depicted in Figure 1, which graphs the video quality scores for one piece of content transcoded using three distinct profiles. While the profile with the highest bit rate (profile 1) maintains a distinctly high score, the medium and low bit rate profiles (2 and 3, respectively) produce very similar results. In this case, dynamic profile selection could have been deployed to transcode the file for profile 3 at a lower bit rate, saving both processing and network resources.

With a dynamic profile selection process, the transcoder uses a two-pass video processing architecture to first examine the source video then transcode it using the individual profile attributes that best align with both the device and the content characteristics. By examining the content before making the profile decision, the transcoder can consistently deliver a video file that meets but does not overly exceed or fall below a set threshold for ABR video quality. This can be done in real time and with a small amount of processing delay.

Dynamic profile selection uses the characteristics of content to determine the perfect mix of all available processing attributes, and in doing so creates versions of content that more precisely match the quality thresholds that service providers define. This means virtually every content version consumes fewer processing resources, takes up less storage space, and requires less precious network bandwidth to deliver. By utilizing these resources more efficiently, service providers can achieve significant cost savings while scaling their multiscreen service offerings to support more content, and more users.

Eliminating Redundant Processing with Cooperative Transcoding

In order to create video files based on each profile, transcoders perform several functions, including decoding, de-interlacing, motion estimation, prediction, de-blocking and entropy coding. Based on today's transcoding architecture, each video source file must be fully processed through of these all stages in order to create the many ABR content versions required. But the results of the decoding, de-interlacing and motion estimation functions are often the same for multiple output streams, presenting opportunities to perform these functions once for many content versions.

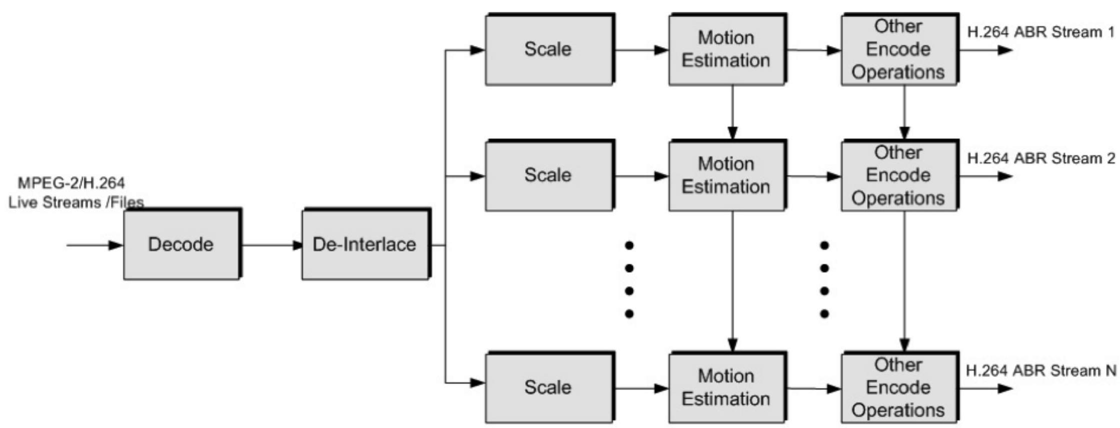

Figure 2: Achieving Efficiencies with Cooperative Transcoding

With cooperative transcoding, instead of conducting every function for every profile for a given piece of content, the transcoder leverages processing modules that are shared between multiple output streams to eliminate redundant processing and save substantial processing resources. This is shown in Figure 2, where the decoding, de-interlacing, motion estimation and other encoding operations for an MPEG-2/H.264 file are shared among ABR streams. According to a recent ARRIS proof of concept demonstration, this type of cooperative transcoding method can reduce transcoder processing requirements by up to 70% for ABR streams at different bitrates, and approximately 30% for ABR streams at different resolutions. These savings present a key opportunity for savings as service providers increase both the volume of content and number of subscribers for multiscreen services.

Conclusion

Together, the efficiencies of cooperative transcoding and dynamic profile selection can generate significant savings in the capital and operational costs of multiscreen content delivery. By creating content profiles dynamically in real-time rather than statically in advance, transcoders may soon be able to custom create a unique profile for each piece of content that delivers the precise attributes of image quality needed, without over or under processing. By trimming even a fraction of unnecessary bits from each content version, service providers can better align their transcoding, storage and networking costs to the multiscreen market opportunity. In addition, by leveraging cooperative transcoding techniques, service providers can achieve even better efficiencies as they roll out multiscreen services on a wide scale.

Together, these two innovations may play a key role in the future of multiscreen video delivery. By making the process of creating, storing and delivering content more efficient, dynamic profile selection and cooperative transcoding may unlock a future where every piece of content, whether linear, stored, popular or long-tail can be delivered to every user – anytime, and anywhere.

Related Reading

- Transcoding Choices for a Multiscreen World – This paper explores the applications for home- and network-based transcoding, and previews some of the innovations that are emerging to help providers transcode their content more efficiently and effectively in the multiscreen world.

- Efficient Content Processing for Adaptive Video Delivery – This paper provides an in-depth overview of two emerging technologies, dynamic profile selection and cooperative transcoding, along with experimental data demonstrating their potential for substantially reducing content processing requirements for multiscreen video delivery.

- Improving Adaptive Video Delivery through Active Management– This paper looks at how service providers can use adaptive delivery technology to deploy unified video processing workflows, which they can use to manage large-scale video delivery over unmanaged networks.